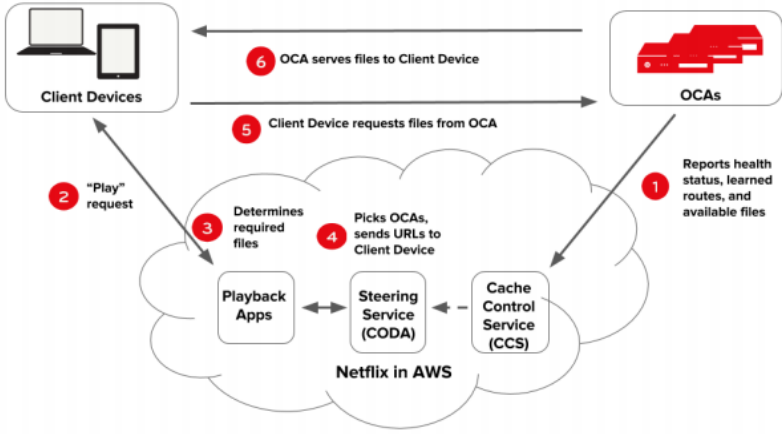


Claim Chart for U.S. Patent 8,495,167 - “Data communications networks, systems, methods and apparatus”

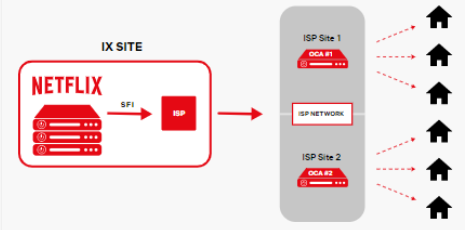
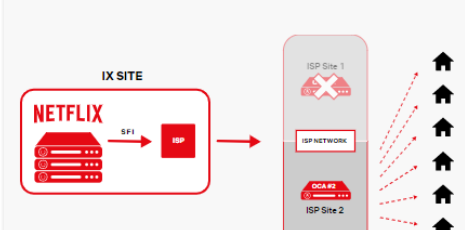
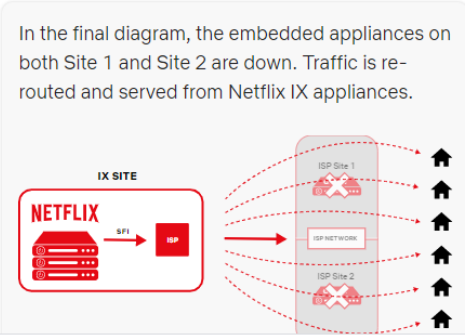
Claim Portion	‘167 Patent	Netflix
Claim 1		
[1a]	A data communication network comprising: a plurality of terminals; and	<p>Netflix uses a system called Open Connect to deliver Netflix TV shows and movies to members world-wide.</p> <p>The building blocks of Open Connect are our suite of purpose-built server appliances, called Open Connect Appliances (OCAs). <i>See</i> Open Connect Overview, p. 2. These are deployed directly inside ISP networks. Netflix provides the server hardware.</p>  <pre> graph TD subgraph Client_Devices [Client Devices] CD[Client Devices] end subgraph Netflix_in_AWS [Netflix in AWS] subgraph AWS_Cloud ([AWS Cloud]) direction TB CCS[Cache Control Service CCS] CODA[Steering Service CODA] PA[Playback Apps] CCS <--> CODA CODA <--> PA end end subgraph OCAs [OCAs] OCA[OCAs] end CD -- "2 'Play' request" --> PA PA -- "3 Determines required files" --> CODA CODA -- "4 Picks OCAs, sends URLs to Client Device" --> CD OCA -- "1 Reports health status, learned routes, and available files" --> CCS CD -- "5 Client Device requests files from OCA" --> OCA OCA -- "6 OCA serves files to Client Device" --> CD </pre>

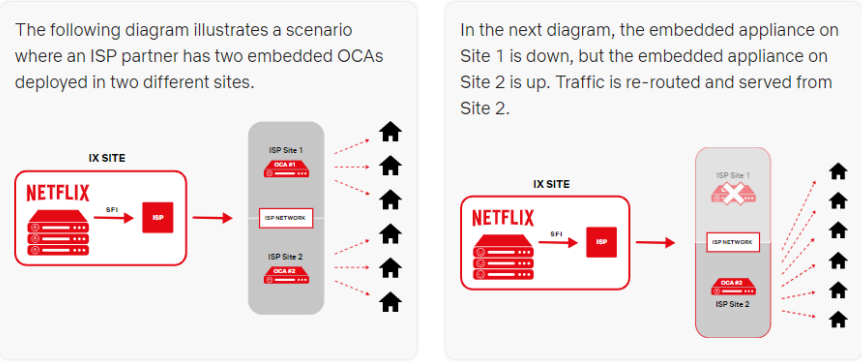
		<p>Open Connect Appliances can be embedded in your ISP network. Embedded OCAs have the same capabilities as the OCAs that we use in our 60+ global data centers, and they are provided to qualifying ISP partners at no charge. Each embedded OCA deployment will offload a substantial amount of Netflix content traffic from peering or transport circuits. Multiple physical deployments can be distributed or clustered on a geographic or network basis to maximize local offload.</p> <p>Source: https://openconnect.netflix.com/en/#sample-architectures</p>
[1b]	a main server adapted to manage selective retrieval of data from a first server by at least one target terminal selected from said plurality of terminals, said main server being distinct from said first server; and	<p>Netflix runs the operation of Open Connect from an application in AWS (main server). When a Netflix user requests playback of a TV show or movie, the AWS application selectively retrieves data from the user device to send content through an internet service provider (ISP). <i>See</i> Open Connect Overview, p. 4-5.</p>

		<ol style="list-style-type: none"> 1. OCA's periodically report health, routes they have learned, and content (file) availability to the cache control services in AWS. 2. A user on a client device requests playback of a title (TV show or movie) from the Netflix application in AWS. 3. The playback application services in AWS check user authorization and licensing, then determine which specific files are required to handle the playback request - taking individual client characteristics and current network conditions into account. <ol style="list-style-type: none"> 4. The steering service in AWS uses the information stored by the cache control service to pick OCAs that the requested files should be served from, generates URLs for these OCAs, and hands the URLs over to the playback application services. 5. The playback application services hand over URLs of the appropriate OCAs to the client device, and the OCA begins to serve the requested files.
[1c]	a network information database containing terminal performance information, wherein	<p>We constantly measure and analyze [OCA] performance and augment capacity as requirements evolve. <i>See</i> Open Connect Overview, p. 3.</p> <p>All OCA deployments are constantly monitored by the Open Connect Operations team to ensure reliability and efficiency. We troubleshoot and proactively fix most issues remotely with minimal input required from our ISP partners. <i>See</i> Open Connect Overview, p. 5.</p>

		<p>Additionally, OCAs periodically report health. <i>Id.</i> at 4.</p> <p>Monitoring, Maintenance, and Updates</p> <p>All of our OCA deployments, whether in IXPs or embedded in ISP networks, are constantly monitored by the Open Connect Operations team to ensure reliability and efficiency. We troubleshoot and proactively fix most issues remotely with minimal input required from our ISP partners. If partners wish to monitor their own embedded OCAs' status and performance, we provide a Partner Portal where they can do so. If hardware performance degrades to the point where a server is no longer functioning in the range of our quality standards, we simply replace it - at no cost to our partners.</p>
[1d]	at least two of said terminals are adapted to act as relay servers for serving data retrieved from said first server to at least one target terminal; and wherein	<p>URLs of appropriate OCAs (relay servers) are given to the client device. The OCAs serve the requested files (Netflix content delivered through ISP) to the client device.</p> <pre> graph TD subgraph Netflix_in_AWS [Netflix in AWS] CCS[Cache Control Service CCS] CODA[Steering Service CODA] PA[Playback Apps] CCS <--> CODA CODA <--> PA end CD[Client Devices] OCA[OCAs] OCA -- 1 Reports health status, learned routes, and available files --> CCS CCS -- 4 Picks OCAs, sends URLs to Client Device --> CD CD -- 5 Client Device requests files from OCA --> OCA OCA -- 6 OCA serves files to Client Device --> CD CD -- 2 "Play" request --> PA PA -- 3 Determines required files --> CODA </pre> <p>The diagram illustrates the Netflix Open Connect architecture. It shows the interaction between Client Devices, OCAs (Open Connect Appliances), and the Netflix in AWS cloud services. The process is numbered 1 through 6:</p> <ol style="list-style-type: none"> OCAs report health status, learned routes, and available files to the Cache Control Service (CCS) in the Netflix in AWS cloud. Client Devices send a "Play" request to the Playback Apps in the Netflix in AWS cloud. Playback Apps determine required files and send the request to the Steering Service (CODA) in the Netflix in AWS cloud. Steering Service (CODA) picks OCAs and sends URLs to the Client Device. Client Device requests files from the OCA. OCA serves files to the Client Device.

[1e]	<p>the main server is adapted to send transport requests direct to at least one first target terminal on the basis of said terminal performance information, and wherein the main server is further adapted to monitor response times of terminals in the network and in which terminals are selected to act as relay servers for a particular data transfers on the basis of their relative response times, and the first target terminal is adapted to act as relay server; and</p>	<p>All OCA deployments are constantly monitored to ensure reliability and efficiency. Netflix makes use of non-peak bandwidth to download the vast majority of content updates to the OCAs in network during these configurable time windows. OCAs can also download updates from each other – minimizing significant usage of internet “backbone” capacity during the update cycle.</p> <p>Netflix also employs multiple failover scenarios where data transfer changes depending on response time of terminals. Further, Netflix can serve traffic directly to a user device.</p> <p>Source: https://openconnect.netflix.com/en/#sample-architectures</p>

		<h3>Failover scenarios for embedded deployments</h3> <p>The following diagram illustrates a scenario where an ISP partner has two embedded OCAs deployed in two different sites.</p>  <p>In the next diagram, the embedded appliance on Site 1 is down, but the embedded appliance on Site 2 is up. Traffic is re-routed and served from Site 2.</p>  <p>In the final diagram, the embedded appliances on both Site 1 and Site 2 are down. Traffic is re-routed and served from Netflix IX appliances.</p> 
[1f]	<p>wherein each such transport request includes details of data to be retrieved, the address of the first server from which the data is to be requested by the first target terminal, the addresses of at least one second target terminal to which the data from the first server to be relayed by the first target terminal and an indication of a relative performance</p>	<p>Netflix sends URLs corresponding to the server from which data is requested by the user device along with other data such as specific files, client characteristics, and network conditions. In setups with multiple OCAs embedded in one ISP partner, data can be sent from one OCA to another in the event of an update or when traffic needs to be re-routed and served from one of the two OCAs because of performance of the network such as when an OCA goes down.</p>

	<p>of a further target terminal based on the terminal performance information stored in the network information database; and</p>	<p>Failover scenarios for embedded deployments</p> <p>The following diagram illustrates a scenario where an ISP partner has two embedded OCAs deployed in two different sites.</p>  <p>In the next diagram, the embedded appliance on Site 1 is down, but the embedded appliance on Site 2 is up. Traffic is re-routed and served from Site 2.</p>
[1g]	<p>wherein terminals adapted to act as relay servers are adapted to modify transport requests received from the main server or from other relay servers and to transmit the modified transport request to selected target terminals from a set of target terminals identified in the transport request, wherein the modified transport request further includes addresses of further target terminals for which the recipient of the modified transport request is to act as relay server; and</p>	<p>OCAs can also download updates from each other – minimizing significant usage of internet “backbone” capacity during the update cycle. The OCAs work in a network to distribute updates among each other and to include further OCAs to which updates and content can be sent. <i>See Open Connect Overview, p. 5.</i></p>

[1h]	<p>wherein data to be retrieved by said target terminals are divided into a series of packets for transmission to said target terminals and each of said terminals is adapted to communicate directly with said main server to acknowledge receipt of the last packet of a series routed thereto.</p>	<h2 data-bbox="842 272 1199 326">“What is TCP?</h2> <p data-bbox="842 354 1877 716">TCP is the de-facto transport protocol on the Internet today and one of the core protocols of the Internet Protocol (IP) suite. It guarantees in-order, error-checked, delivery of all content sent from one network device to another. TCP employs retransmissions to ensure that no portion of the content is lost. To that end, TCP breaks content into packets. Each packet has a sequence number that identifies its relative ordering. The sender transmits packets to the receiver and expects acknowledgements for in-order, correctly received, packets. If any packet is detected as lost, it is retransmitted. TCP is responsible for re-arranging the received packets and delivering the content to the application without errors or data gaps.</p> <p data-bbox="842 756 869 777">...</p> <h2 data-bbox="842 808 1549 862">Where does Netflix use TCP?</h2> <p data-bbox="842 889 1814 997">Netflix uses TCP for internet streaming to send packets of data for video. Additionally, Netflix specifically looks at the number of TCP connections to determine internet speeds in accordance with testing of OCAs.</p> <p data-bbox="842 1081 1845 1151"><i>On information and belief, Netflix’s terminals also communicate with the Main server to acknowledge receipt of the data packets.</i></p>
------	---	--

CONNEC

PEERING
WITH
OPEN
CONNECTOPEN
CONNECT
APPLIANCESDEPLOYING
EMBEDDED
APPLIANCESPARTNER
PORTAL

SUPPORT



HOME

[WELCOME TO OPEN CONNECT >](#)[WHAT IS OPEN CONNECT >](#)[HOW TO GET STARTED >](#)[SAMPLE ARCHITECTURES >](#)[WHAT'S HERE? >](#)

Open Connect

HOME

[WELCOME TO OPEN CONNECT >](#)[WHAT IS OPEN CONNECT >](#)[HOW TO GET STARTED >](#)[SAMPLE ARCHITECTURES >](#)[WHAT'S HERE? >](#)

Welcome to Open Connect

The goal of the Netflix Open Connect program is to provide our millions of Netflix subscribers the highest-quality viewing experience possible. We achieve this goal by partnering with Internet Service Providers (ISPs) to deliver our content more efficiently. We partner with over a thousand ISPs to localize substantial amounts of traffic with Open Connect Appliance embedded deployments, and we have an open peering policy at our [interconnection locations](#). If you are an ISP with a substantial amount of Netflix traffic, review this information to learn more about the program.

For more information about Open Connect, see:

- [Overview of Open Connect \(PDF - English only\)](#)
- [Open Connect blog post](#)

☐ Key links on this site:

CONNEC

PEERING
WITH
OPEN
CONNECTOPEN
CONNECT
APPLIANCESDEPLOYING
EMBEDDED
APPLIANCESPARTNER
PORTAL

SUPPORT



Fill out the [appliance request form](#) if you are interested in embedded appliance solutions.

What is Open Connect

The Netflix Open Connect program provides opportunities for ISP partners to improve their customers' Netflix user experience by localizing Netflix traffic and minimizing the delivery of traffic that is served over a transit provider.

There are two main components of the program, which are architected in partnership with ISPs to provide maximum benefit in each individual situation: embedded Open Connect Appliances and settlement-free interconnection (SFI).

Embedded Open Connect Appliances (OCAs)

Open Connect Appliances can be embedded in your ISP network. Embedded OCAs have the same capabilities as the OCAs that we use in our 60+ global data centers, and they are provided to qualifying ISP partners at no charge. Each embedded OCA deployment will offload a substantial amount of Netflix content traffic from peering or transport circuits. Multiple physical deployments can be distributed or clustered on a geographic or network basis to maximize local offload.

Netflix provides:

- ☐ • Network architecture and technical turn-up expertise
- Ongoing monitoring and issue resolution

CONNEC

PEERING
WITH
OPEN
CONNECTOPEN
CONNECT
APPLIANCESDEPLOYING
EMBEDDED
APPLIANCESPARTNER
PORTAL

SUPPORT

▼ 文A

(

The [OCA Deployment Guide](#) provides deployment details. If you are interested in OCA deployment at your ISP, please fill out our [OCA request form](#).

Settlement-free interconnection (SFI)

Connect via direct Private Network Interconnect (PNI) or IXP-based SFI peering to Netflix Open Connect Appliances in our data centers. Peering alone can be beneficial.

If you deploy embedded OCAs, you also set up SFI peering for additional resiliency and to enable nightly content fills and updates.

Netflix has the ability to interconnect at a number of global data center facilities and public Internet Exchange fabrics as listed on our [Peering Locations](#) page. We openly peer with any network at IXP locations where we are mutually present and we consider private interconnection as appropriate. If you are interested in interconnection, please review the information on the [Peering Locations](#) page.

ISPs who do not currently participate in public peering might want to consider that a single IX port can support multiple peering sessions, providing direct access to various content, cloud, and network providers. In addition to Netflix, many large organizations such as Akamai, Amazon, Facebook, and Google/YouTube widely participate in public peering and combine to deliver a substantial percentage of traffic to a typical ISP.

From a connectivity standpoint, IX ports can be reached locally in a data center or via transport. We recommend <http://peeringdb.com> as a detailed source of information that can help you find an IX that best meets your needs.



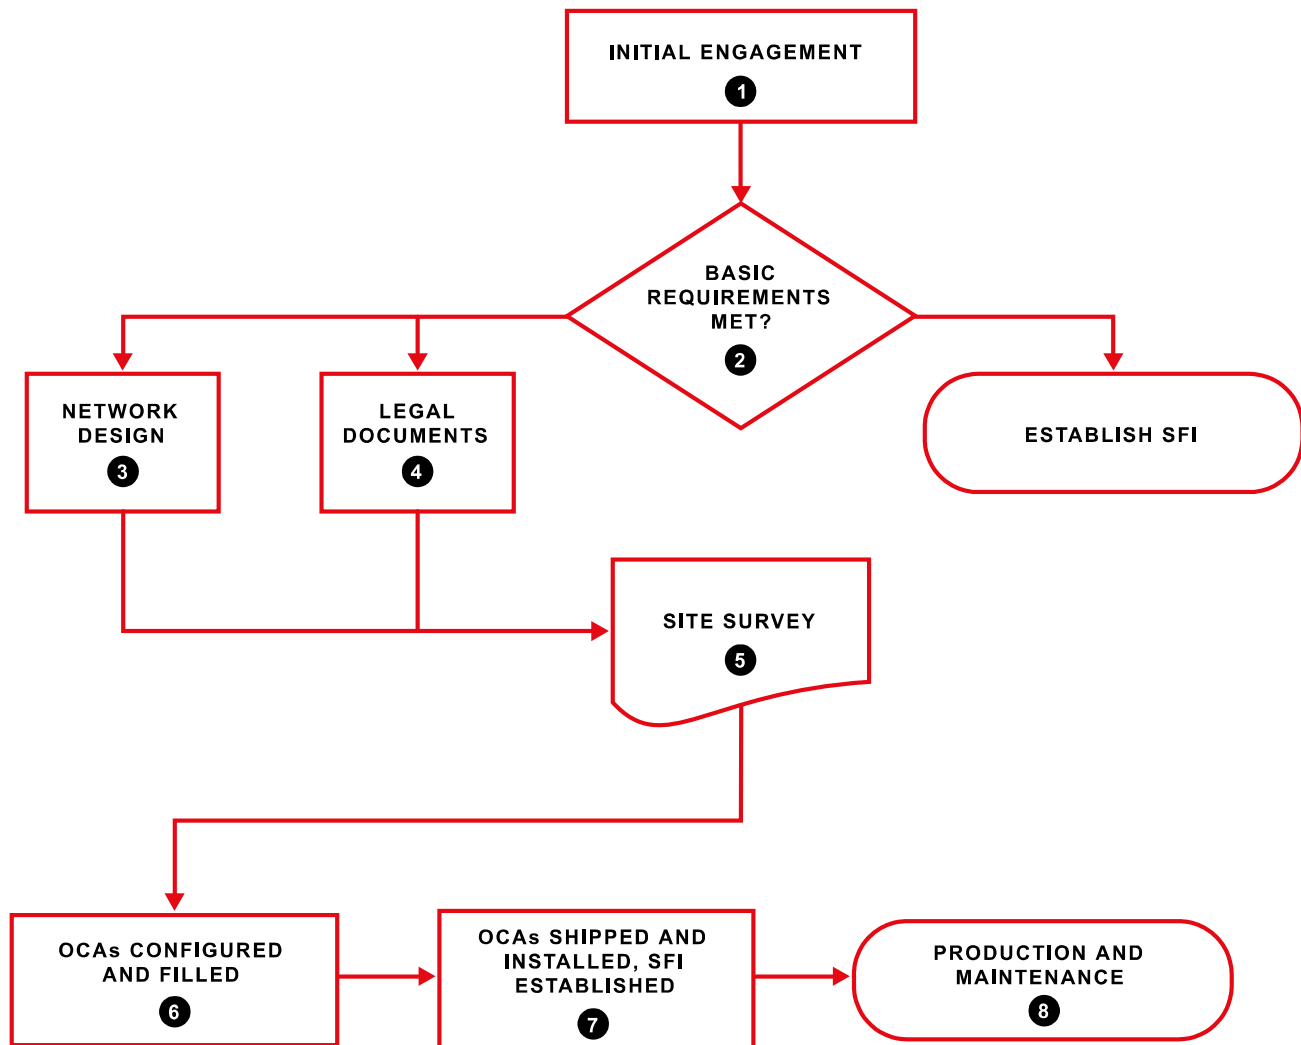
CONNEC

PEERING
WITH
OPEN
CONNECTOPEN
CONNECT
APPLIANCESDEPLOYING
EMBEDDED
APPLIANCESPARTNER
PORTAL

SUPPORT

A

(



1 To begin the engagement process, you complete one of the following actions:

- If you are engaging with Netflix for the first time, please submit an [online Open Connect Appliance request](#), and a Netflix Open Connect Partner Engagement Manager (PEM) will contact you.
- If you already have an established relationship with a Netflix PEM and you are interested in modifying an existing configuration, contact your PEM directly.

2 Netflix works with you to determine whether you meet the [basic requirements for deploying embedded OCAs](#). If embedded OCAs are not warranted or cannot be deployed for some reason, SFI peering alone can be established.



CONNEC

PEERING
WITH
OPEN
CONNECTOPEN
CONNECT
APPLIANCESDEPLOYING
EMBEDDED
APPLIANCESPARTNER
PORTAL

SUPPORT

文A


- 7 Netflix ships you the configured and filled appliance(s) for installation in your data center. SFI is also established to enable optimal resiliency and fill/updates. Shipping typically takes approximately **1 week**, depending on your geographical location.
- 8 After the appliance is enabled into production, Netflix performs ongoing **health monitoring and maintenance** to ensure smooth operation of the appliance.

Sample Architectures

The following examples illustrate how Open Connect deployments can be configured for resiliency and maximum benefit to the customer.

Basic Architectures

The following diagram shows an example of an OCA that is embedded in a partner network, in conjunction with SFI peering which is used to provide additional resiliency and to enable nightly content **fill and updates**.



In contrast, the next diagram shows an example of SFI (peering) *without* the deployment of embedded OCAs in the partner network. In this scenario, traffic is delivered to end users via SFI from Netflix appliances that are located in local IXPs, to avoid both the cost and congestion that is associated with transit.



CONNEC

PEERING
WITH
OPEN
CONNECTOPEN
CONNECT
APPLIANCESDEPLOYING
EMBEDDED
APPLIANCESPARTNER
PORTAL


SUPPORT

文A

(

Failover scenarios for embedded deployments

The following diagram illustrates a scenario where an ISP partner has two embedded OCAs deployed in two different sites.



In the next diagram, the embedded appliance on Site 1 is down, but the embedded appliance on Site 2 is up. Traffic is re-routed and served from Site 2.

CONNEC

PEERING
WITH
OPEN
CONNECT


OPEN
CONNECT
APPLIANCES

DEPLOYING
EMBEDDED
APPLIANCES

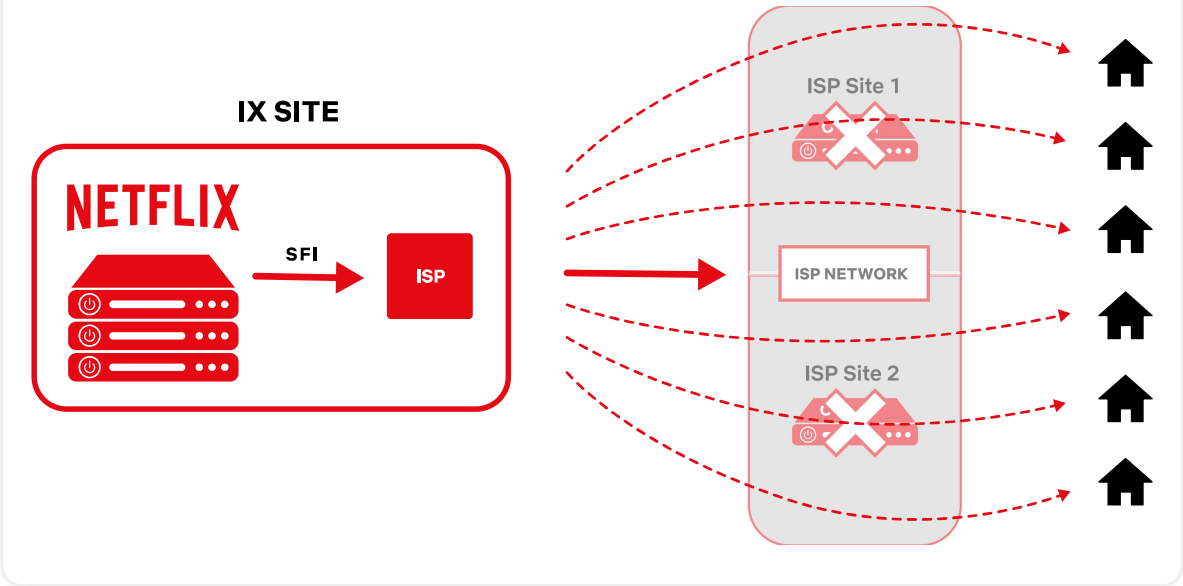
PARTNER
PORTAL

SUPPORT

文A



In the final diagram, the embedded appliances on both Site 1 and Site 2 are down. Traffic is re-routed and served from Netflix IX appliances.




CONNEXPEERING
WITH
OPEN
CONNECTOPEN
CONNECT
APPLIANCESDEPLOYING
EMBEDDED
APPLIANCESPARTNER
PORTAL

SUPPORT



Peering With Open Connect



Open Connect Appliances



Deploying Embedded Appliances



Partner Portal



Frequently Asked Questions



Jobs



Building fast.com



Netflix Technology Blog

Follow

Aug 9, 2016 • 7 min read

On [our company blog in May](#), we introduced [fast.com](#), our new internet speed test. The idea behind fast.com is to provide a quick and simple way for any internet user to test their current internet speed, whether they are a Netflix member or not. Since fast.com was released, millions of internet users around the world have run the test. We have seen a lot of interest in the site and questions about how it works. This blog will give a high-level overview of how we handled some of the [challenges inherent with measuring internet speeds](#) and the technology behind fast.com.

But first, some news — we are happy to announce a new FAST mobile app, available now for Android or Apple mobile devices. Get the free app from the [Apple App Store](#) or [Google Play](#).

Design goals

When designing the user experience for the fast.com application, we had several important goals in mind:

- Provide accurate, consistent results that reflect users' real-life internet use case
- Load and run as quickly as possible
- Provide simple results that are easy to understand
- Work on most devices from the browser without requiring installation of a separate application

We wanted to make sure that fast.com could be easily used and understood by the majority of internet users, without requiring them to have any prior knowledge of

computer networking, command line tools, and the like.

Technical goals

There are various ways to go about measuring internet speed and many variables that can impact any given measurement, some of which are not under our control. For example — configuration of the user's local or home network, device or router performance, other users on the network, TCP or network configuration on the device. However, we thought carefully about the variables that are under our control and how they would further our overall goal of a simple but meaningful test.

Variables that are under our control, and which can influence the results of the test, include things like:

- Server location
- Load on the server
- Number of TCP connections used
- Size and type of download content used
- Methodology used to aggregate measurements

One major advantage we have is our Open Connect CDN, a globally-distributed network of servers (Open Connect Appliances or OCAs) that store and serve Netflix content to our members — representing as much as 35% of last-mile internet peak traffic in some regions. Using our own production servers to test internet speed helps to ensure that the test is a good representation of the performance that can be achieved during a real-life user scenario.

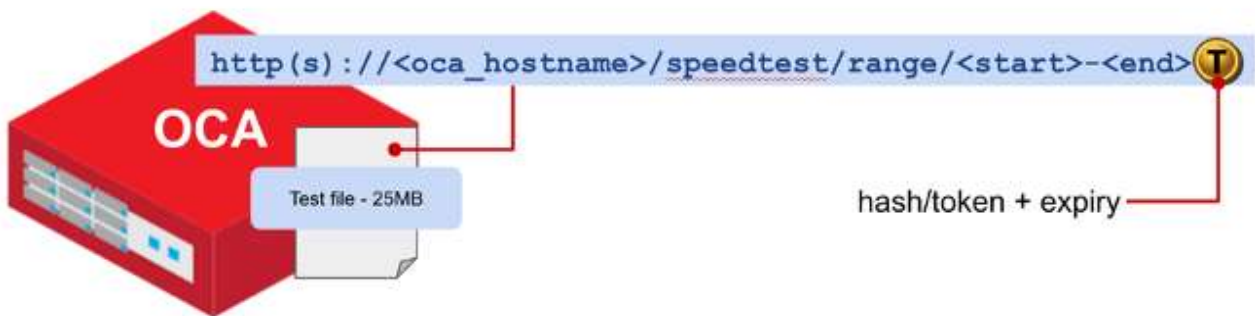
In pursuit of the design goal of simplicity, we deliberately chose to measure only download speed, measuring how fast data travels from server to consumer when they are performing activities such as viewing web pages or streaming video. Downloads represent the majority of activity for most internet consumers.

We also decided on the following high-level technical approaches:

- To open several connections for the test, varying the number depending on the network conditions
- To run the test on several of our wide network of Netflix production OCAs, but only on servers that have enough capacity to serve test traffic while simultaneously operating within acceptable parameters to deliver optimal video quality to members
- To measure long running sessions — eliminating connection setup and ramp up time and short term variability from the result
- To dynamically determine when to end the test so that the final results are quick, stable, and accurate
- To run the test using HTTPS, supporting IPv4 and IPv6

Architecture

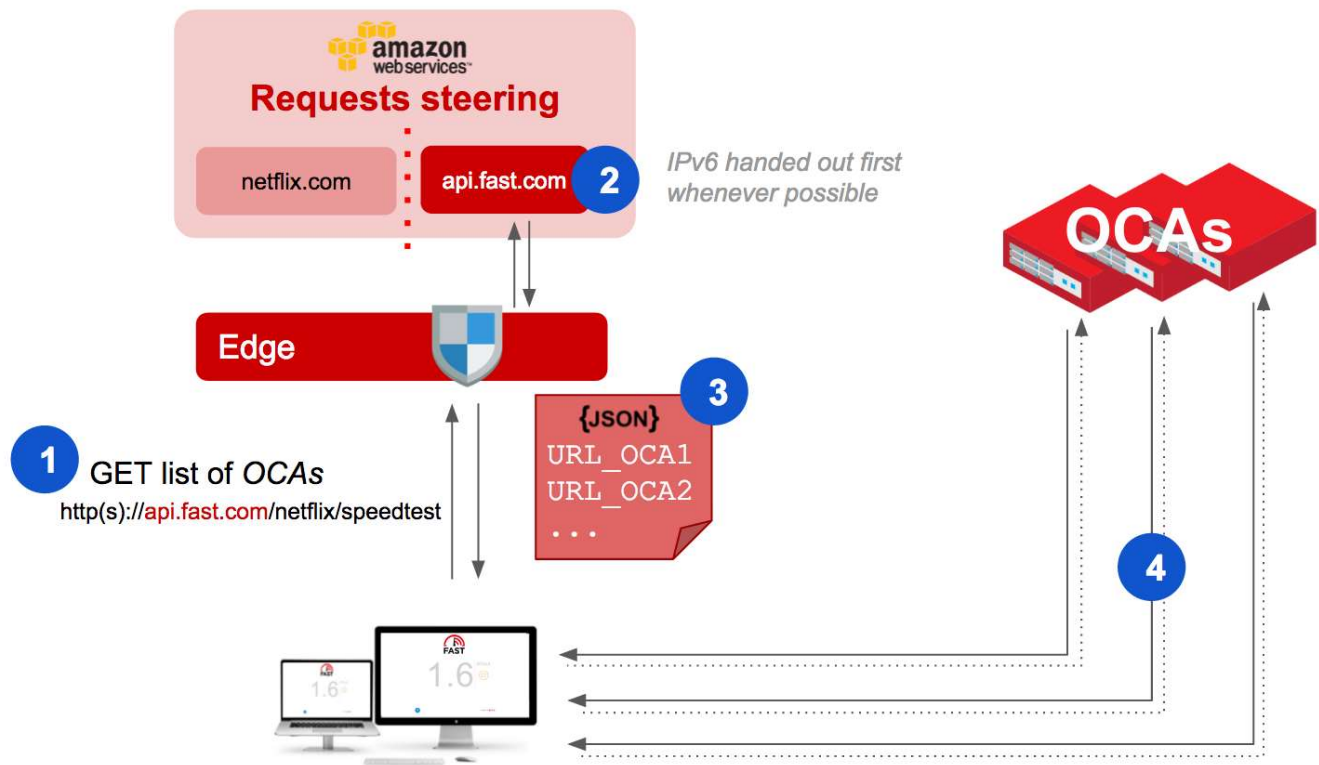
As mentioned above, fast.com downloads test files from our distributed network of Open Connect Appliances (OCAs). Each OCA server provides an endpoint with a 25MB video file. The endpoint supports a range parameter that allows requests for between a 1 byte to a 25MB chunk of content.



In order to steer a user to an OCA server, fast.com provides an endpoint that returns a list of several URLs for different OCAs that are best suited to run the test. To determine the list, the endpoint uses logic that is similar to the logic that is used to steer netflix.com video delivery. The OCAs that are returned are chosen based on:

- Network distance

- Traffic load for each OCA, which indicates overall server health
- Network structure — each OCA in the list belongs to a different cluster



As soon as the fast.com client receives the URLs, the test begins to run.

Estimating network speed

The test engine uses heuristics to:

- Strip off measurements that are collected during connection setup/ramp up
- Aggregate the rest of the collected measurements
- Decide how many parallel connections to use during the test
- Try to separate processing overhead from network time — because fast.com runs in the browser, it has limited visibility into timing of network events like DNS resolution time, processing of packets on the client side and latency to test server

- Make a decision about when the client has collected enough measurements to confidently present the final network speed estimate

We exclude initial connection ramp up, but we do take into account any performance drops during the test. Network performance drops might indicate a lossy network, congested link, or faulty router — therefore, excluding these drops from the test result would not correctly reflect issues experienced by users while they are consuming content from the internet.

Number of connections

Depending on network throughput, the fast.com client runs the test using a variable number of parallel connections. For low throughput networks, running more connections might result in each connection competing for very limited bandwidth, causing more timeouts and resulting in a longer and less accurate test.

When the bandwidth is high enough, however, running more parallel connections helps to saturate the network link faster and reduce test time. For very high throughput connections, especially in situations with higher latency, one connection and a 25MB file might not be enough to reach maximum speeds, so multiple connections are necessary.

Size of downloads

For each connection, the fast.com client selects the size of the chunk of the 25MB file that it wants to download. In situations where the network layer supports periodical progress events, it makes sense to request the whole file and estimate network speed using download progress counters. In cases where the download progress event is not available, the client will gradually increase payload size during the test to perform multiple downloads and get a sufficient number of samples.

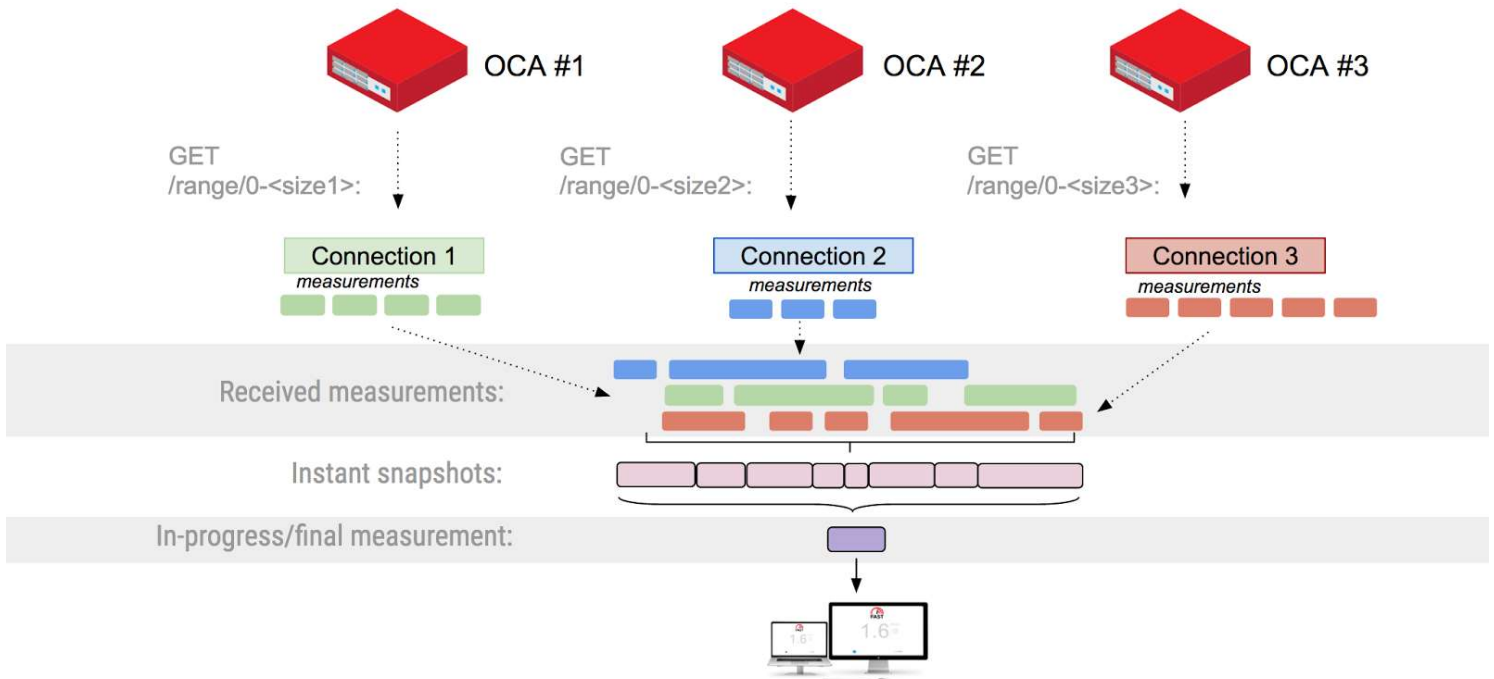
Computing the results

After the download measurements are collected, the client combines the downloaded content across all connections and keeps the snapshot speed.

The ‘instant’ network measurements are then passed to the results aggregation module. The aggregation module makes sure that:

- We exclude initial connection ramp up

- We take the rest and compute rolling average of the other measurements




One of the primary challenges for the fast.com client is determining when the estimated speed measurements are ready to be presented as a final estimate. Due to the various environments and conditions that the fast.com test can be run under, the test duration needs to be dynamic.


For stable low latency connections, we quickly see growth to full network speeds:



Higher latency connections take much longer to ramp up to full network speed:



Lossy or congested connections show significant variations in instant speed, but these instant variations get smoothed out over time. It is also harder to correctly identify the moment when connections have ramped up to full speed.



In all cases, after initial ramp up measurements are excluded, the ‘stop’ detection module monitors how the aggregated network speed is changing and makes a decision about whether the estimate is stable or if more time is needed for the test. After the results are stable, they are presented as a final estimate to the user.

Conclusion and Next Steps

We continue to monitor, test, and perfect fast.com, always with the goal of giving consumers the simplest and most accurate tool possible to measure their current internet performance. We plan to share updates and more details about this exciting tool in future posts.

— by Sergey Fedorov and Ellen Livengood

Originally published at techblog.netflix.com on August 9, 2016.

[Network Speed](#) [Internet Speed](#) [Isps](#)

[About](#) [Write](#) [Help](#) [Legal](#)

Get the Medium app





Open Connect Overview

What is Netflix Open Connect?

Open Connect is the name of the global network that is responsible for delivering Netflix TV shows and movies to our members world-wide. This type of network is typically referred to as a “Content Delivery Network” or “CDN” because its job is to deliver internet-based content (via HTTP/HTTPS) efficiently by bringing the content that people watch close to where they’re watching it. The Open Connect network shares some characteristics with other CDNs, but also has some important differences.

Netflix began the Open Connect initiative in 2011, as a response to the ever-increasing scale of Netflix streaming. We started the program for two reasons:

- 1) As Netflix grew to be a significant portion of overall traffic on consumer Internet Service Provider (ISP) networks, it became important to be able to work with those ISPs in a direct and collaborative way.
- 2) Creating a content delivery solution customized for Netflix allowed us to design a proactive, directed caching solution that is much more efficient than the standard demand-driven CDN solution, reducing the overall demand on upstream network capacity by several orders of magnitude.

Several years in, we are constantly learning - adjusting and evolving the program to ensure that our members continue to have a consistently high quality video experience from wherever they are in the world.

The overall mission of the program is to enable ISPs to provide a great Netflix experience for our mutual customers. We further this goal by localizing Netflix traffic as close as possible to our members, limiting the network and geographical distances that our video bits must travel during playback. This of course benefits Netflix members, but it also benefits ISPs and internet users in general. In short, we invest in efficiency

innovations and increasing the capacity of the internet to support playback requests for Netflix content - so that others don't have to.


Open Connect Appliances

The building blocks of Open Connect are our suite of purpose-built server appliances, called Open Connect Appliances (OCAs). These appliances store encoded video/image files and serve these files via HTTP/HTTPS to client devices (for example: set top boxes, mobile devices, or smart TVs). OCAs have the sole responsibility of delivering playable bits to client devices as fast as possible.


As with all facets of the Open Connect program, appliance design continues to expand and improve over time to keep up with our current and future needs. We are involved in the development of all layers of the software stack, and we make both the hardware design and the software available via open source for others to benefit from.

Our global network of thousands of OCAs are deployed in two ways:

1. We install OCAs within internet exchange points (referred to as IXs or IXPs) in significant Netflix markets throughout the world. These OCAs are interconnected with mutually-present ISPs via settlement-free public or private peering (SFI). Peering alone can be very beneficial to our ISP partners.



2. We provide OCAs free of charge to qualifying ISPs. These OCAs, with the same capabilities as the OCAs that are in the IXPs, are deployed directly inside ISP networks. We provide the server hardware and the ISPs provide power, space, and connectivity. ISPs directly control which of their customers are routed to their embedded OCAs. ISP partners with embedded OCAs also use peering for resiliency and to enable efficient nightly fill and updates, described later.



Each individual OCA deployment site is custom tailored by the Open Connect team based on local network characteristics and other key capacity planning factors. After we deploy OCAs to a site, we constantly measure and analyze their performance and augment capacity as requirements evolve.

Interaction with Client Devices and Netflix AWS Services

OCAs do not store client data (for example - viewing history, DRM info, or member data). Essentially, OCA servers only do the following two things:

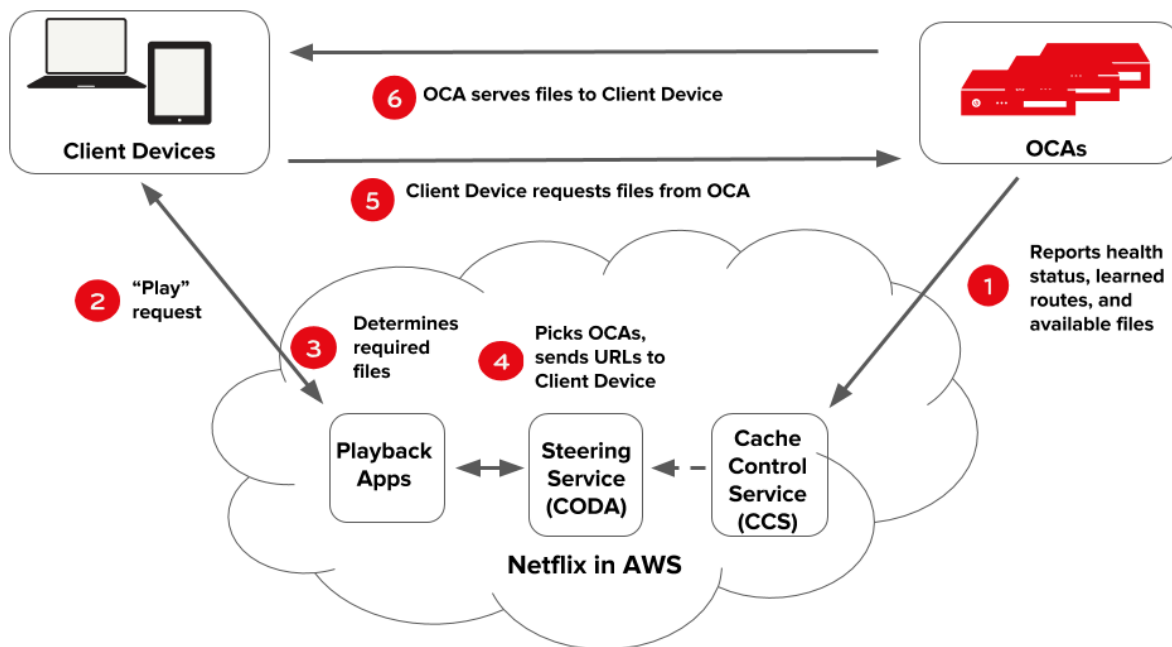
- Report their status to the Open Connect control plane services in Amazon Web Services (AWS). For example, they report health metrics, BGP routes they have learned from the BGP peer (router or switch) they have a configured BGP session with, and what files they have stored on disk.
- Serve content via HTTP/HTTPS when it is requested by a client device.

The control plane services in AWS take the data that the OCAs report and use it to steer clients via URL to the most optimal OCAs given their file availability, health, and network

proximity to the client. The control plane services also control fill behavior (adding new files to OCAs nightly), compute optimal behavior for such things as file storage/hashing, and handle the storage and interpretation of relevant telemetry about the playback experience.

Open Connect also partners extensively with Netflix client teams to ensure that the content that is served by the OCAs is optimized dynamically by each client device based on its specific needs and the current network conditions.

The following diagram illustrates how the playback process works:



1. OCA's periodically report health, routes they have learned, and content (file) availability to the cache control services in AWS.
2. A user on a client device requests playback of a title (TV show or movie) from the Netflix application in AWS.
3. The playback application services in AWS check user authorization and licensing, then determine which specific files are required to handle the playback request - taking individual client characteristics and current network conditions into account.

4. The steering service in AWS uses the information stored by the cache control service to pick OCAs that the requested files should be served from, generates URLs for these OCAs, and hands the URLs over to the playback application services.
5. The playback application services hand over URLs of the appropriate OCAs to the client device, and the OCA begins to serve the requested files.

Monitoring, Maintenance, and Updates

All of our OCA deployments, whether in IXPs or embedded in ISP networks, are constantly monitored by the Open Connect Operations team to ensure reliability and efficiency. We troubleshoot and proactively fix most issues remotely with minimal input required from our ISP partners. If partners wish to monitor their own embedded OCAs' status and performance, we provide a Partner Portal where they can do so. If hardware performance degrades to the point where a server is no longer functioning in the range of our quality standards, we simply replace it - at no cost to our partners.

We have the somewhat unique benefit of being able to deploy the majority of our content and software updates proactively during off-peak fill windows. Because we can predict with high accuracy what our members will watch and what time of day they will watch it, we make use of non-peak bandwidth to download the vast majority of content updates to the OCAs in our network during these configurable time windows. OCAs can also download updates from each other - minimizing significant usage of internet "backbone" capacity during the update cycle.

We use ever-evolving popularity algorithms and storage techniques in the control plane services to ensure that our content is distributed in ways that maximize offload efficiency and quality of experience, while minimizing churn in the form of updates to the content that is stored on the appliances.

More Information

For more information about Open Connect, see:

- [The Open Connect website](#)
- [Open Connect articles on the Netflix Tech Blog](#)
- [Open Connect articles on the Netflix Media Center](#)

Copyright © 2016, 2019 by Netflix, Inc. All rights reserved. No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or by any information storage or retrieval system, without express permission from Netflix, Inc.